U.S. Department of Energy's
Office of Science

Connecting Scientists

…to collaborators, to facilities, to datastores, to computational resources, to….

Energy Sciences Network (ESnet)

September 24, 2004
HEPAP Meeting

Mary Anne Scott
Program Manager
Advanced Scientific Computing Research
Mathematical, Information, and Computational Research Division
scott@er.doe.gov

What is ESnet?

- Mission:
  - Provide, interoperable, effective and reliable communications infrastructure and leading-edge network services that support missions of the Department of Energy, especially the Office of Science
- Vision:
  - Provide seamless and ubiquitous access, via shared collaborative information and computational environments, to the facilities, data, and colleagues needed to accomplish their goals.
- Role:
  - A component of the Office of Science infrastructure critical to the success of its research programs (funded through ASCR/MICS and managed and operated by ESnet staff at LBNL).

Essentially all of the national data traffic supporting US science is carried by two networks—ESnet and Internet-2/Abilene (which plays a similar role for the university community)

What is ESnet's user base?

- Betweeen 10,000 and 100,000 researchers in the US (guesstimate)
- Mainly Office of Science programs—ASCR, BER, BES, FES, HEP, NP
- Also traffic for NNSA and others
- All the US national labs
- Hundreds of universities
- Hundreds of foreign institutions

Characteristics of the user base

- Many casual users
- Users from science disciplines that span SC interests
- Users concerned with any data intensive and computationally intensive tasks
- Collaborators distributed geographically, small to large groups

Who Generates Traffic, and Where Does it Go?

ESnet Inter-Sector Traffic Summary, Jan 2003 / Feb 2004 (1.7X overall traffic increase, 1.9X OSC increase)

(the international traffic is increasing due to BABAR at SLAC and the LHC tier 1 centers at FNAL and BNL)

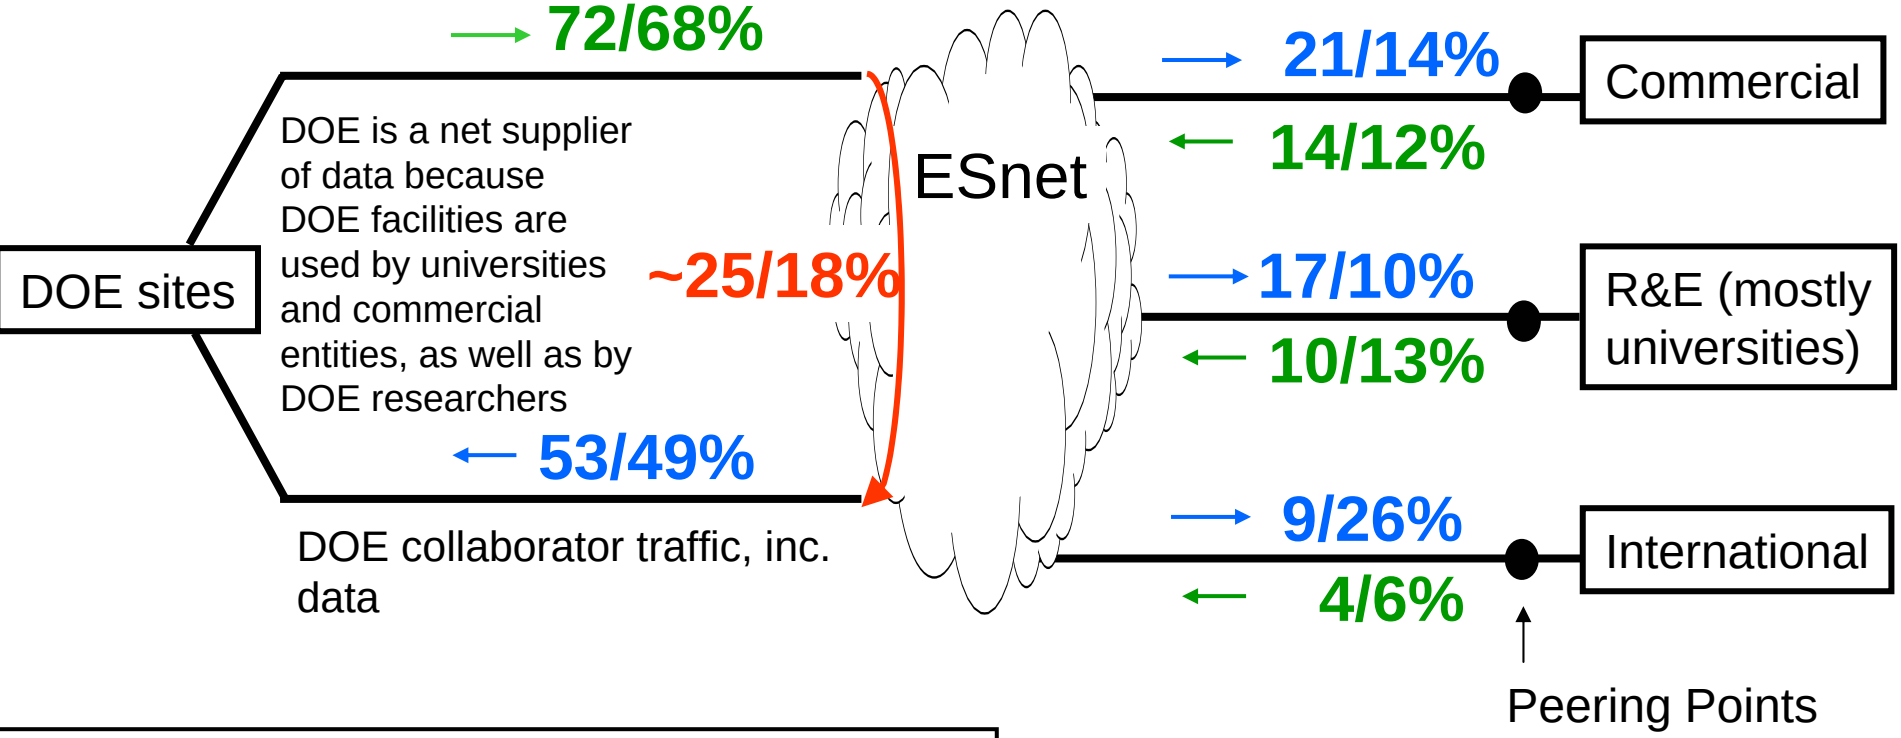

Note that more that 90% of the ESnet traffic is OSC traffic

ESnet Appropriate Use Policy (AUP)

All ESnet traffic must originate and/or terminate on an ESnet an site (no transit traffic is allowed)

Traffic coming into ESnet = Green
Traffic leaving ESnet = Blue
Traffic between sites
% = of total ingress or egress traffic

ESnet is Engineered to Move a Lot of Data

ESnet Top 20 Data Flows, 24 hr. avg., 2004-04-20

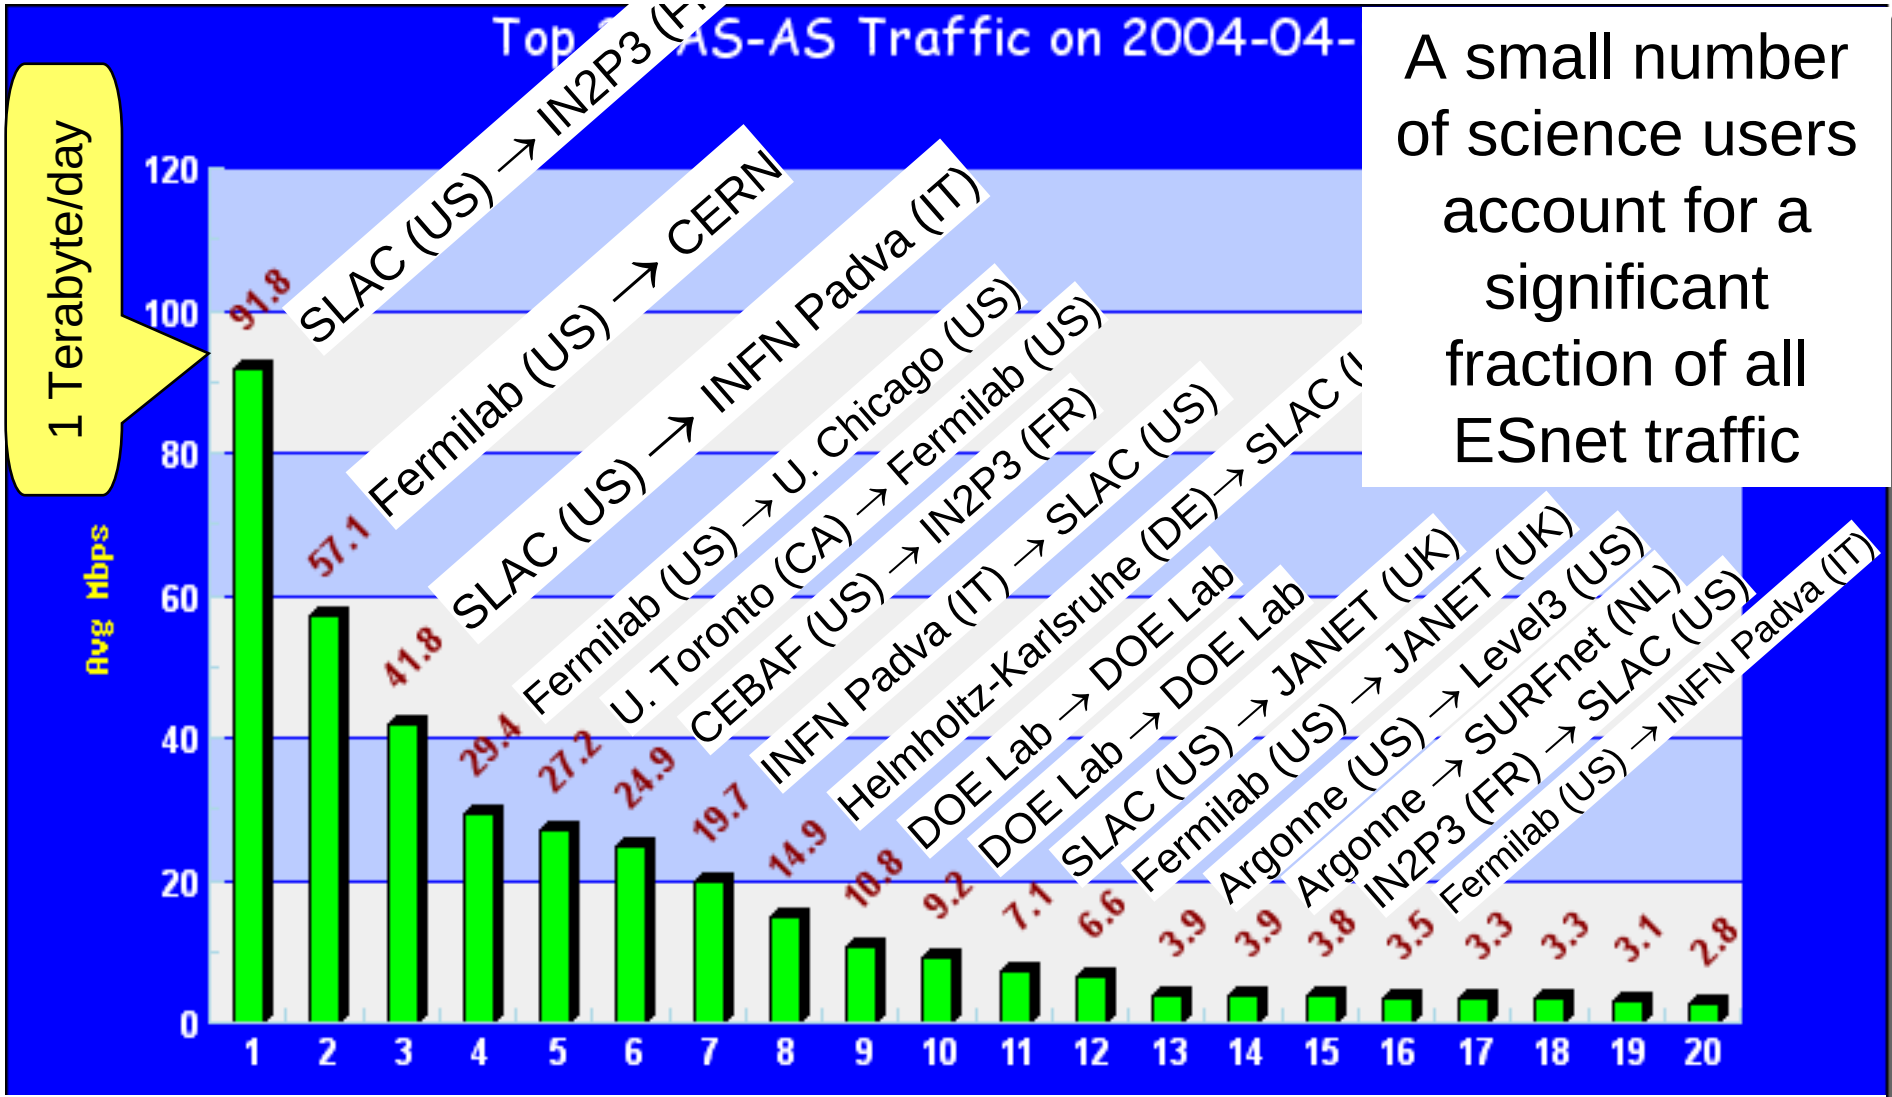

A small number of science users account for a significant fraction of all ESnet traffic

➢Since BaBar data analysis started, the top 20 ESnet flows have consistently accounted for > 50% of ESnet's monthly total traffic (~130 of 250 TBy/mo)

- The traffic is not transient: Daily and weekly averages are about the same.
- SLAC is a prototype for what will happen when Climate, Fusion, SNS, Astrophysics, etc., start to ramp up the next generation science

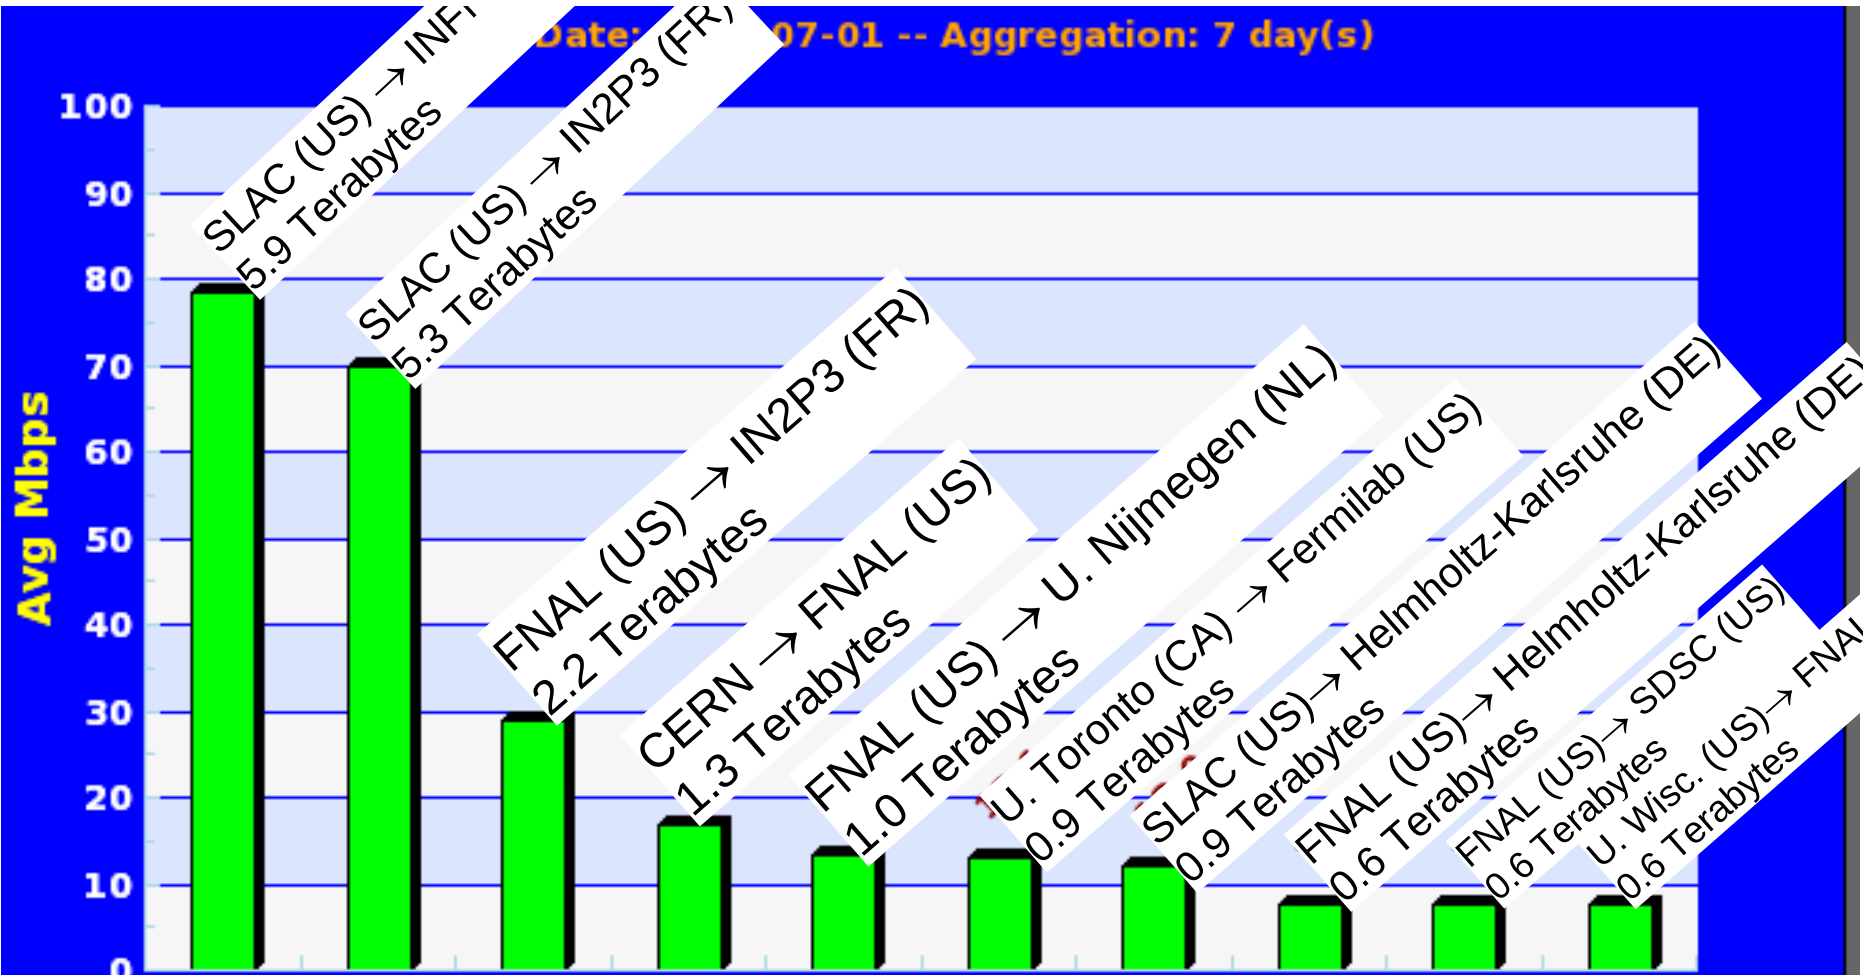

ESnet Top 10 Data Flows, 1 week avg., 2004-07-01

Why is ESnet important?

- Enables thousands of DOE, university and industry scientists and collaborators worldwide to make effective use of unique DOE research facilities and computing resources independent of time and geographic location
  - Direct connections to all major DOE sites
  - Access to the global Internet (managing 150,000 routes at 10 commercial peering points)
  - User demand has grown by a factor of more than 10,000 since its inception in the mid 1990's—a 100 percent increase every year since 1990
- Capabilities not available through commercial networks
  - Architected to move huge amounts of data between a small number of sites
  - High bandwidth peering to provide access to US, European, Asia-Pacific, and other research and education networks.

Objective: *Support scientific research* by providing seamless and ubiquitous access to the facilities, data, and colleagues

Science Mission Critical Infrastructure

- ESnet is a visible and critical piece of DOE science infrastructure
  - If ESnet fails, tens of thousands of DOE and University users know it within minutes if not seconds
- Requires high reliability and high operational security in the
  - network operations, and
  - ESnet infrastructure support – the systems that support the operation and management of the network and services
    - Secure and redundant mail and Web systems are central to the operation and security of ESnet
      - trouble tickets are by email
      - engineering communication by email
      - engineering database interface is via Web
    - Secure network access to Hub equipment
    - Backup secure telephony access to all routers
    - 24x7 help desk (joint w/ NERSC) and 24x7 on-call network engineers

ESnet Connects DOE Facilities and Collaborators

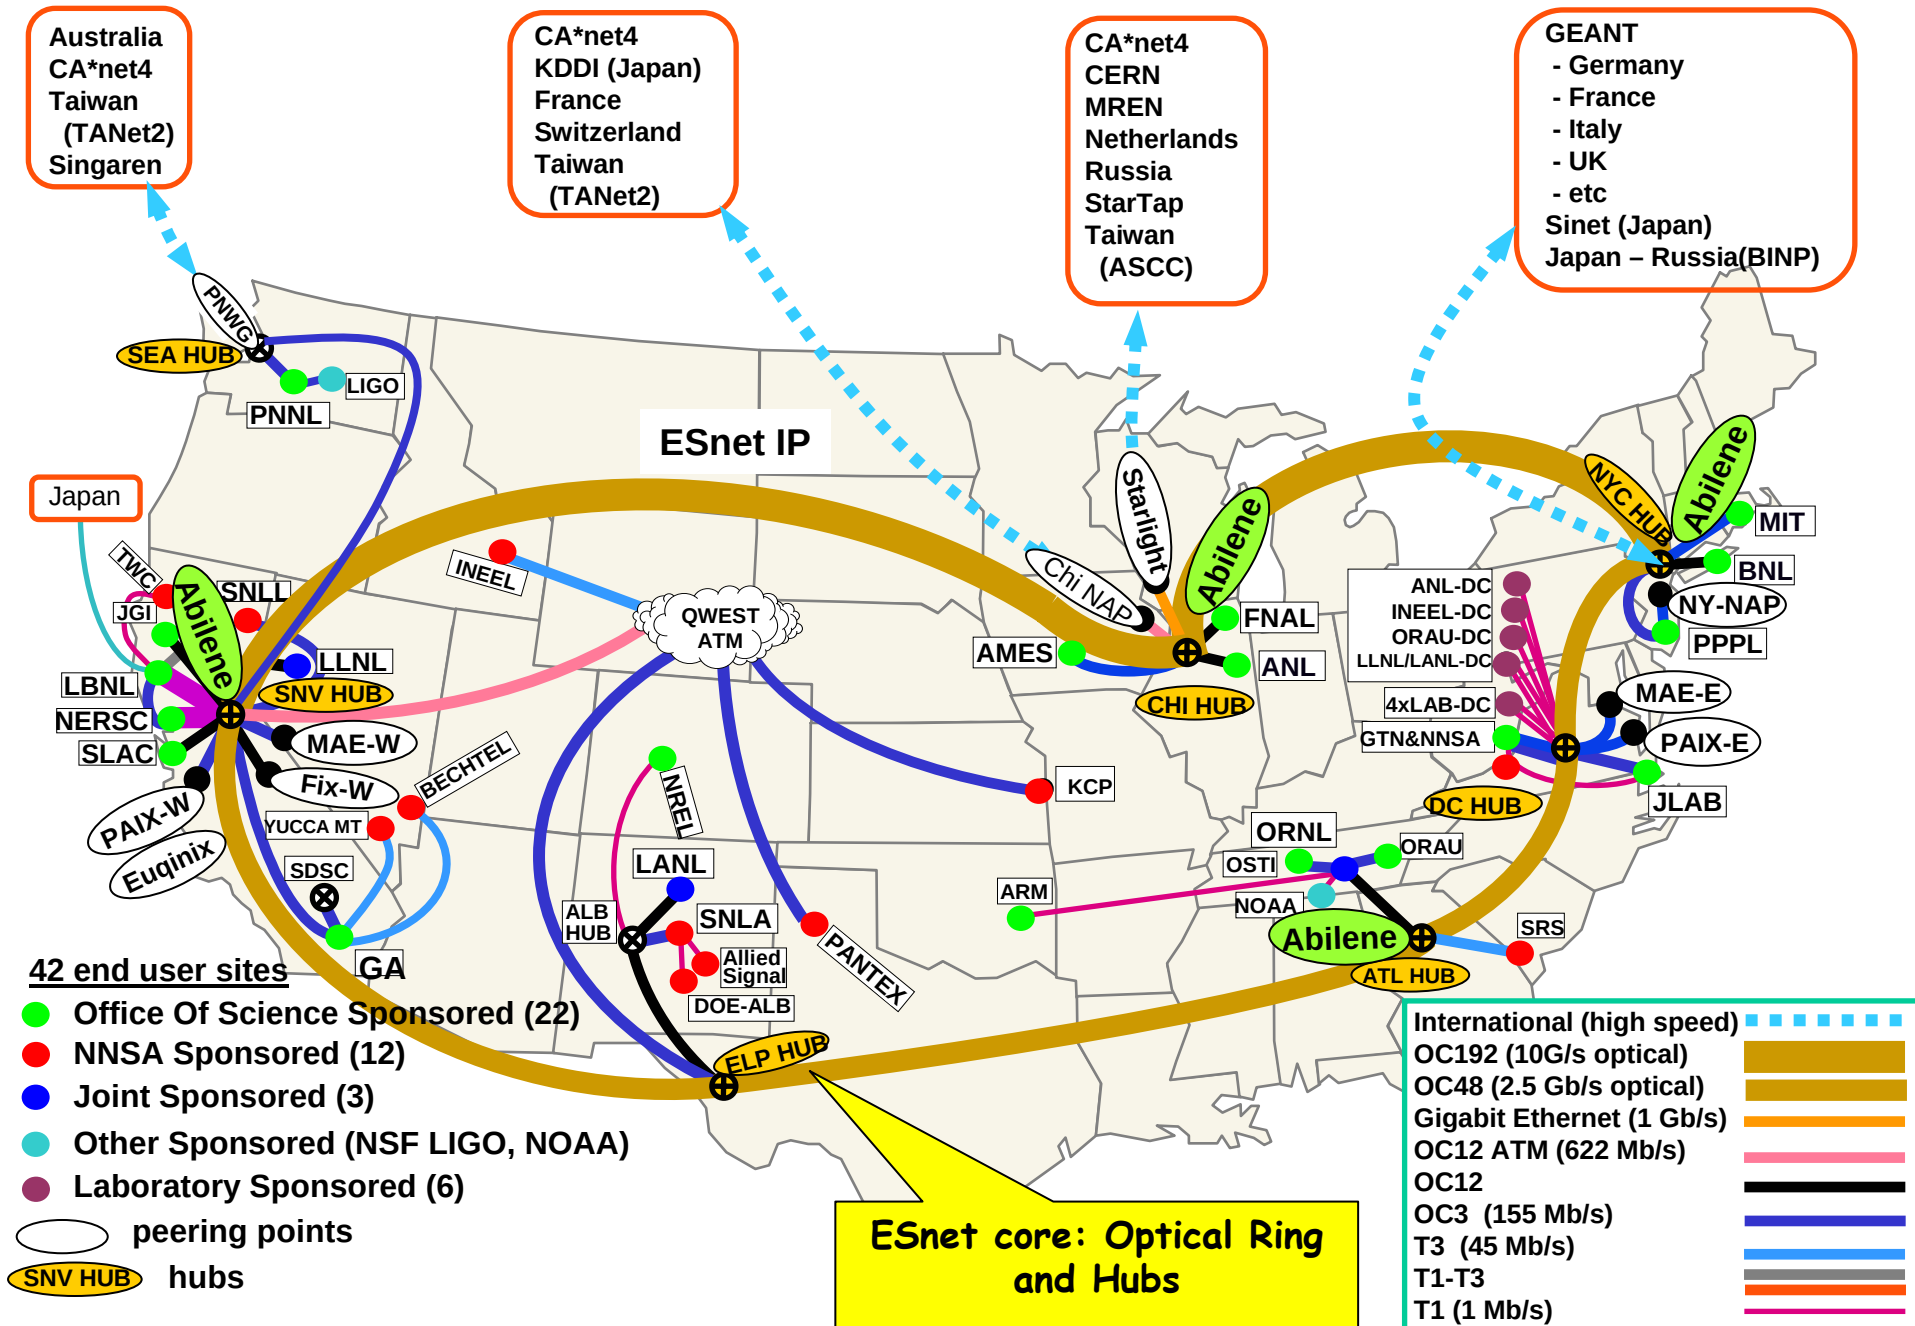

ESnet's Peering Infrastructure
Connects the DOE Community With its Collaborators

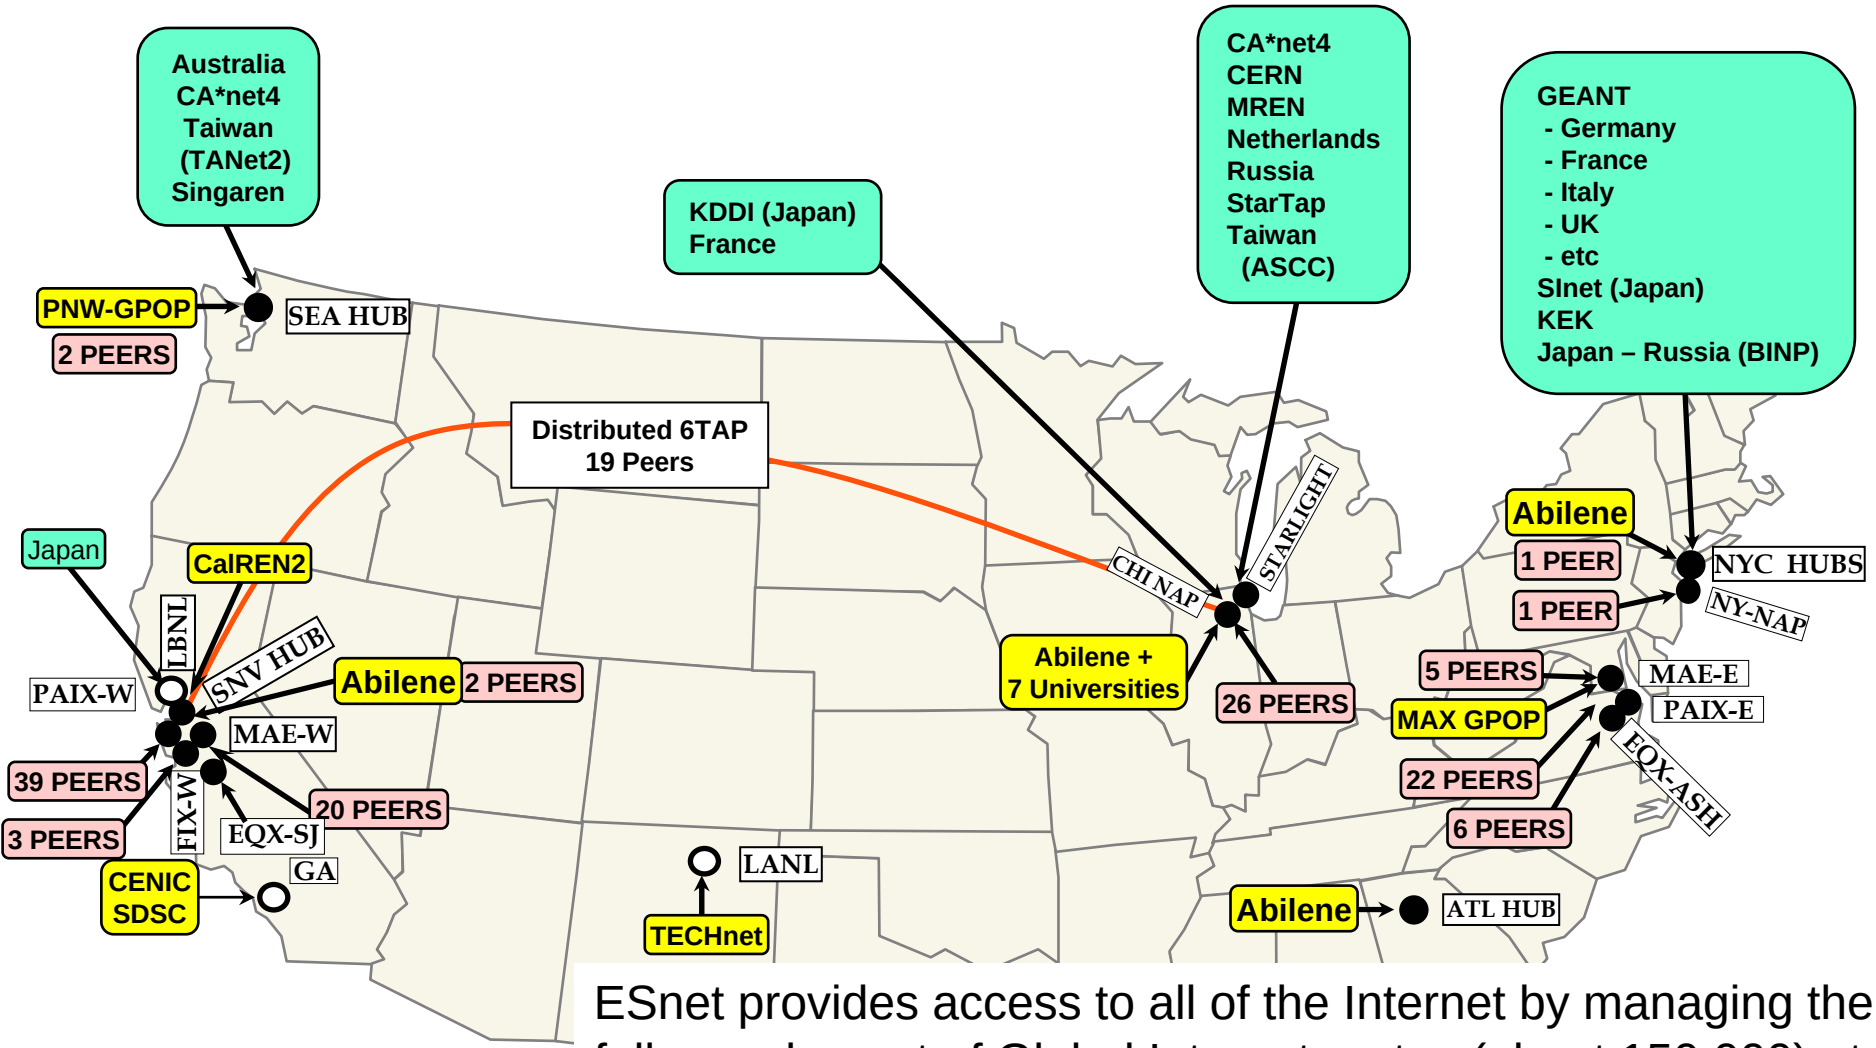

ESnet Peering (connections to other networks)

- University
- International
- Commercial

ESnet provides access to all of the Internet by managing the full complement of Global Internet routes (about 150,000) at 10 general/commercial peering points + high-speed peerings w/ Abilene and the international R&E networks. *This is a lot of work, and is very visible, but provides full access for DOE.*

How is ESnet Managed?

- A community endeavor
  - Strategic guidance from the OSC programs
    - Energy Science Network Steering Committee (ESSC)
      - High Energy Physics represented by Larry Price, ANL; Richard Mount, SLAC
    - Network operation is a shared activity with the community
    - ESnet Site Coordinators Committee
    - Ensures the right operational "sociology" for success
- Complex and specialized – both in the network engineering and the network management – in order to provide its services to the laboratories in an integrated support environment
- Extremely reliable in several dimensions
- ➢ Taken together these points make ESnet a unique facility supporting DOE science that is quite different from a commercial ISP or University network

U.S. Department of Energy

Office of Science

ESnet WAN Security and Cyberattack Defense

- Cyber defense is a new dimension of ESnet security
  - Security is now inherently a global problem
  - As the entity with a global view of the network, ESnet has an important role in overall security

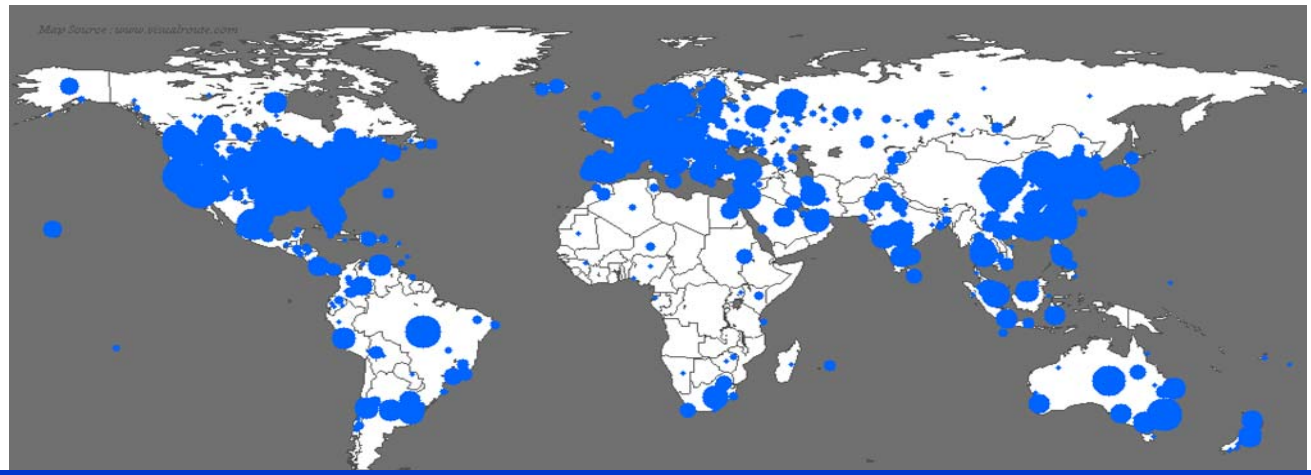

30 minutes after the Sapphire/Slammer worm was released, 75,000 hosts running Microsoft's SQL Server (port 1434) were infected.

("The Spread of the Sapphire/Slammer Worm," David Moore (CAIDA & UCSD CSE), Vern Paxson (ICIR & LBNL), Stefan Savage (UCSD CSE), Colleen Shannon (CAIDA), Stuart Staniford (Silicon Defense), Nicholas Weaver (Silicon Defense & UC Berkeley EECS) http://www.cs.berkeley.edu/~nweaver/sapphire) Jan., 2003

U.S. Department of Energy

Office of Science

ESnet and Cyberattack Defense

Sapphire/Slammer worm infection hits creating almost a full Gb/s (1000 megabit/sec.) traffic spike on the ESnet backbone

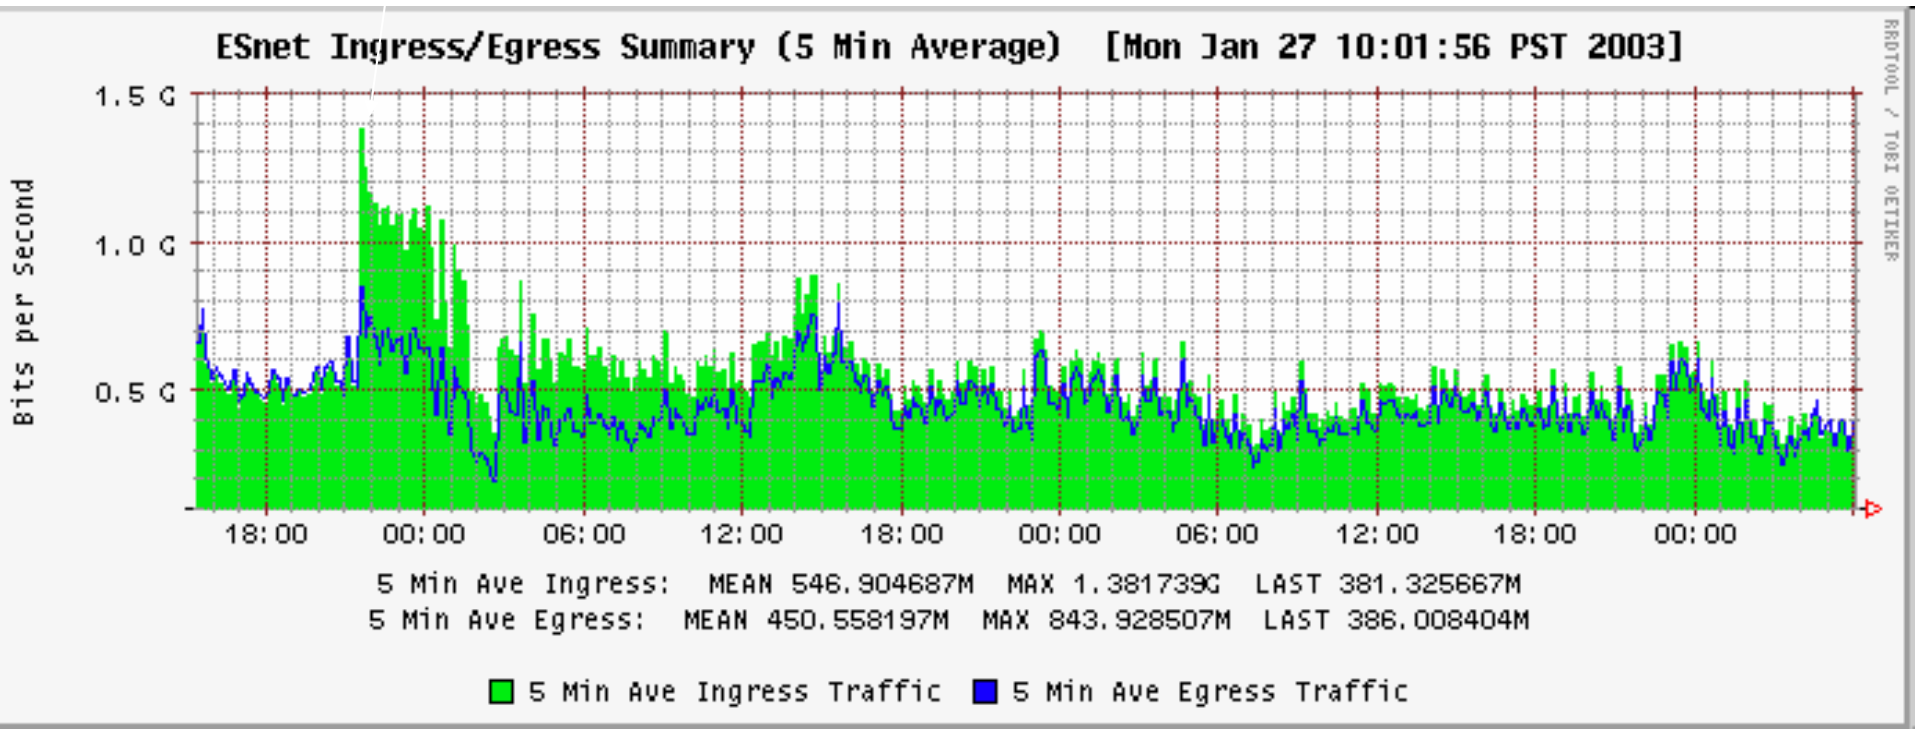

U.S. Department of Energy

Office of Science

Planning Workshops

- **High Performance Network Planning Workshop, August 2002**
  http://www.doecollaboratory.org/meetings/hpnpw
- **DOE Workshop on Ultra High-Speed Transport Protocols and Network Provisioning for Large-Scale Science Applications, April 2003**
  http://www.csm.ornl.gov/ghpn/wk2003
- **Science Case for Large Scale Simulation, June 2003**
  http://www.pnl.gov/scales/
- **DOE Science Networking Roadmap Meeting, June 2003**
  http://www.es.net/hypertext/welcome/pr/Roadmap/index.html
- **Workshop on the Road Map for the Revitalization of High End Computing, June 2003**
  http://www.cra.org/Activities/workshops/nitrd
  http://www.sc.doe.gov/ascr/20040510_hecrtf.pdf **(public report)**
- **ASCR Strategic Planning Workshop, July 2003**
  http://www.fp-mcs.anl.gov/ascr-july03spw
- **Planning Workshops-Office of Science Data-Management Strategy, March & May 2004**
  - http://www-conf.slac.stanford.edu/dmw2004 **(report coming soon)**

What's a possible future?

VISION –

- A seamless, high-performance network infrastructure in which science applications and advanced facilities are "n-way" interconnected to terascale computing, petascale storage, and high-end visualization capabilities.
- This advanced network facilitates collaborations among researchers and interactions between researchers and experimental and computational resources,
- Science, especially large-scale science, moves to a new regime that eliminates isolation, discourages redundancy, and promotes rapid scientific progress through the interplay of theory, simulation, and experiment.

U.S. Department of Energy

Office of Science

Network and Middleware Needs of DOE Science

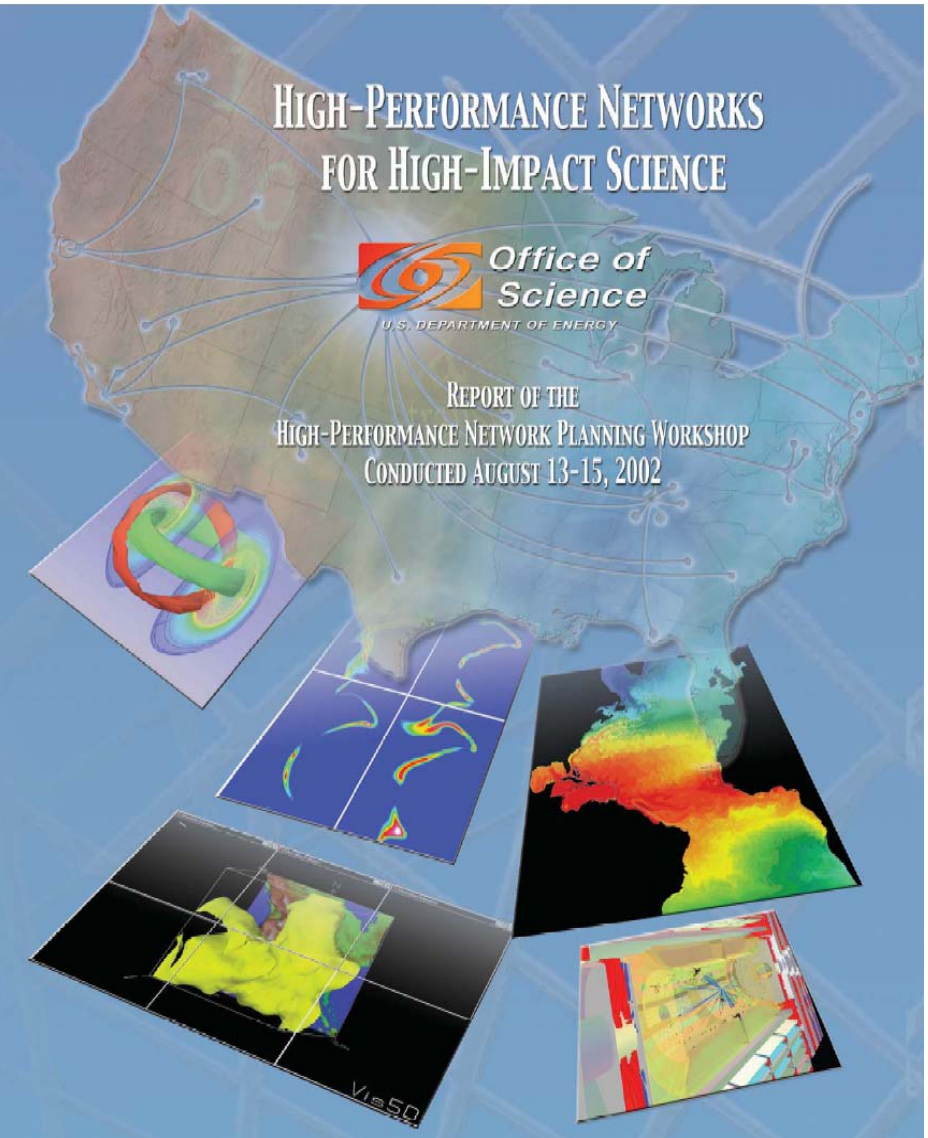

- Focused on science requirements driving Advanced network infrastructure
  - Middleware research
  - Network research
  - Network governance model
- Requirements for DOE science developed for a representative cross-section of the OSC scientific community

Evolving Quantitative Science Requirements for Networks

Science Areas	Today *End2End* Throughput	5 years End2End Throughput	5-10 Years End2End Throughput	Remarks
High Energy Physics	0.5 Gb/s	100 Gb/s	**1000 Gb/s**	high bulk throughput
Climate (Data & Computation)	0.5 Gb/s	160-200 Gb/s	**N x 1000 Gb/s**	high bulk throughput
SNS NanoScience	Not yet started	1 Gb/s	**1000 Gb/s + QoS for control channel**	remote control and time critical throughput
Fusion Energy	0.066 Gb/s (500 MB/s burst)	0.198 Gb/s (500MB/ 20 sec. burst)	**N x 1000 Gb/s**	time critical throughput
Astrophysics	0.013 Gb/s (1 TBy/week)	N*N multicast	**1000 Gb/s**	computational steering and collaborations
Genomics Data & Computation	0.091 Gb/s (1 TBy/day)	100s of users	**1000 Gb/s + QoS for control channel**	high throughput and steering

U.S. Department of Energy

Office of Science

New Strategic Directions to Address Needs of DOE Science

Focused on what is needed to achieve the science driven network requirements of the previous workshop

- **THE #1 DRIVER for continuing advancements in networking and middleware** – Petabyte scale experimental and simulation data systems will be increasing to exabyte scale data systems.
  - Bioinformatics, Climate, LHC, etc.
- Computational systems that process or produce data continue to advance with Moore's Law
- …..

Organized by the ESSC

- Workshop Chair, Roy Whitney, JLAB
- Workshop Editors, Roy Whitney, JLAB; Larry Price, ANL

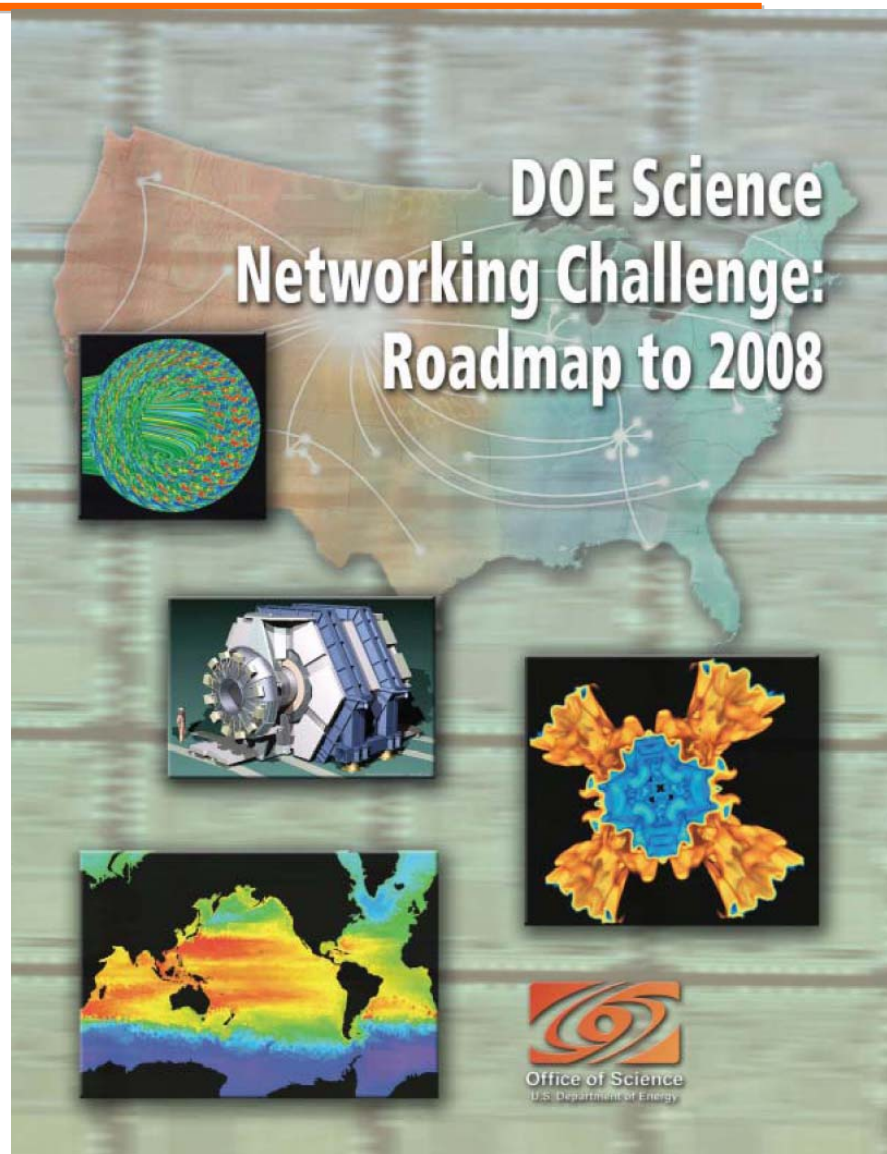

Observed Drivers for ESnet Evolution

- Are we seeing the predictions of two years ago come true?
- Yes!

…what now???

VISION - A scalable, secure, integrated *network environment* for ultra-scale distributed science is being developed to make it possible to combine resources and expertise to address complex questions that no single institution could manage alone.

- Network Strategy

  Production network

  - Base TCP/IP services; +99.9% reliable

  High-impact network

  - Increments of 10 Gbps; switched lambdas (other solutions); 99% reliable

  Research network

  - Interfaces with production, high-impact and other research networks; start electronic and advance towards optical switching; very flexible

- Revisit governance model
  - SC-wide coordination
  - Advisory Committee involvement

Evolution of ESnet

- Upgrading ESnet to accommodate the anticipated increase from the current 100%/yr traffic growth to 300%/yr over the next 5-10 years is priority number 7 out of 20 in DOE's "Facilities for the Future of Science – A Twenty Year Outlook"
- Based on the requirements, ESnet must address

  I. Capable, scalable, and reliable production IP networking
    - University and international collaborator connectivity
    - Scalable, reliable, and high bandwidth site connectivity

  II. Network support of high-impact science
    - provisioned circuits with guaranteed quality of service (e.g. dedicated bandwidth)

  III. Evolution to optical switched networks
    - Partnership with UltraScienceNet
    - Close collaboration with the network R&D community

  IV. Science Services to support Grids, collaboratories, etc

New ESnet Architecture to Accommodate SC

- The future requirements cannot be met with the current, telecom provided, hub and spoke architecture of ESnet

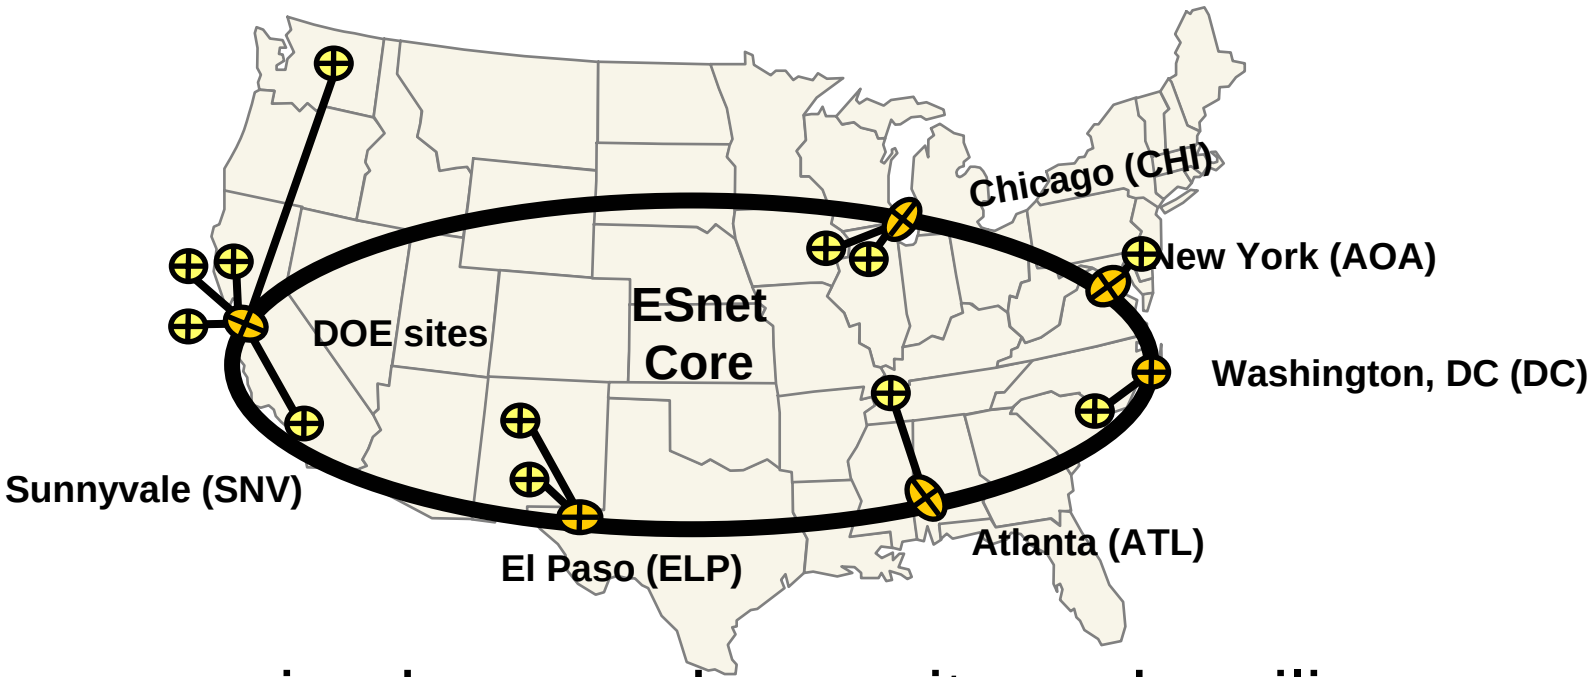

- The core ring has good capacity and resiliency against single point failures, but the point-to-point tail circuits are neither reliable nor scalable to the required bandwidth

- ESnet new architecture goals: full redundant connectivity for every site and high-speed access for every site (at least 10 Gb/s)
- Two part strategy
  1) MAN rings provide dual site connectivity and much higher site bandwidth
  2) A second backbone will provide
     - multiply connected MAN rings for protection against hub failure
     - extra backbone capacity
     - a platform for provisioned, guaranteed bandwidth circuits
     - alternate path for production IP traffic
     - carrier neutral hubs

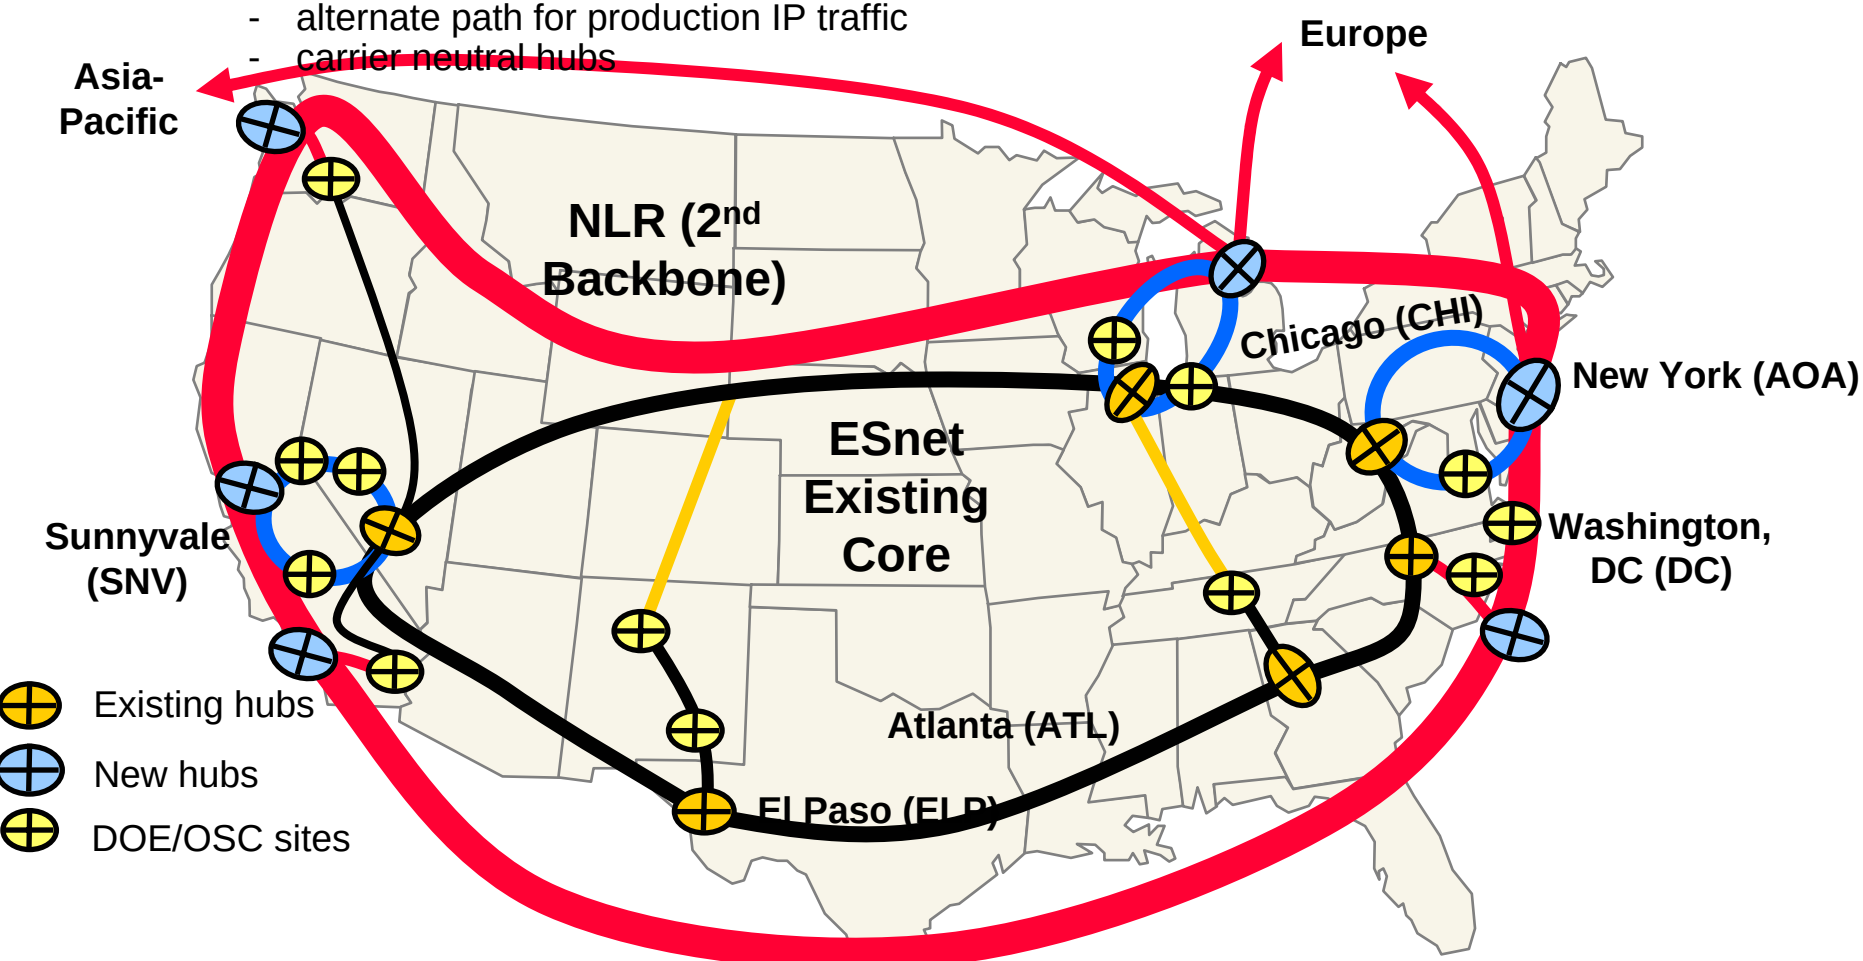

Science Impact from PPDG

- **Robust, sustained, hands-off, production data transfer of terabytes of data has been enabled using GridFTP and SRM implementations.**
  - **D0 moved over 50 TB of event data (20% of the run) to be analyzed off-site from FNAL over the past six months. By using multiple streams in GridFTP, a *factor of 5 improvement in throughput* was made possible.**
- **Grid-based job scheduling and execution based on Condor-G, DAGMan and GRAM:**
  - **Single scientists on PPDG experiments are able to schedule and manage significant numbers of jobs at multiple sites on the grid. This is now done using *a single FTE in cases where 2 or more FTEs* were previously required.**
  - **CMS simulated event data is now entirely Grid based. Jobs submitted using MOP and the CMS Distributed Production Environment based on Condor-G and Globus through use of the Virtual Data Toolkit have generated over 50 million events with an overall *factor of two increased efficiency* than a year ago. Over 75,000 jobs have been run supported by a single FTE.**

Conclusions

- ESnet is an infrastructure that is critical to DOE's science mission
- Focused on the Office of Science Labs, but serves many other parts of DOE
- ESnet is working hard to meet the current and future networking need of DOE mission science in several ways:
  - Evolving a new high speed, high reliability, leveraged architecture
  - Championing several new initiatives which will keep ESnet's contributions relevant to the needs of our community

Where do you come in?

- Early identification of requirements
  - Evolving programs
  - New facilities
- Interaction with HEP representatives on ESSC
- Participation in management activities
- Next ESSC meeting early CY2005 in DC area